ARTist: The Android Runtime Instrumentation and Security Toolkit

Michael Backes[1], Sven Bugiel[2], Oliver Schranz[2], Philipp von Styp-Rekowsky[2]
Sebastian Weisgerber[2]

[1]CISPA, Saarland University & MPI-SWS, Saarland Informatics Campus
[2]CISPA, Saarland University, Saarland Informatics Campus

arXiv:1607.06619v1 [cs.CR] 22 Jul 2016

ABSTRACT

We present *ARTist*, a compiler-based application instrumentation solution for Android. *ARTist* is based on the new ART runtime and the on-device *dex2oat* compiler of Android, which replaced the interpreter-based managed runtime (DVM) from Android version 5 onwards. Since *dex2oat* is yet uncharted, our approach required first and foremost a thorough study of the compiler suite's internals and in particular of the new default compiler backend *Optimizing*. We document the results of this study in this paper to facilitate independent research on this topic and exemplify the viability of *ARTist* by realizing two use cases. Moreover, given that seminal works like TaintDroid hitherto depend on the now abandoned DVM, we conduct a case study on whether taint tracking can be re-instantiated using a compiler-based instrumentation framework. Overall, our results provide compelling arguments for preferring compiler-based instrumentation over alternative bytecode or binary rewriting approaches.

1. INTRODUCTION

Google's Android OS has become a popular subject of the security research community over the last few years. Among the different directions of research on improving Android's security, a dedicated line of work has successfully investigated how instrumentation of the interpreter (i.e., Dalvik virtual machine) can be leveraged for security purposes. This line of work comprises influencing works such as TaintDroid [17] for analyzing privacy relevant data flows within applications, AppFence [22] for protecting the end-users' privacy, Moses [28] for domain isolation, or Spandex [13] for password tracking, just to name a few.

However, with the release of Android 5 Lollipop, Google made a large technological leap by replacing the interpreter-based runtime with an on-device, ahead-of-time compilation of apps to platform specific bytecode that is executed in the new Android runtime (short ART). While this leap did not affect the app developers, it broke legacy compliance of all of the previously mentioned security solutions that rely on instrumentation of the DVM and it restricts them to Android versions prior to Lollipop. In fact, it has left the security research community with two choices for carrying on work that relies on instrumented runtimes: resorting to binary or bytecode rewriting techniques [16, 21] or adapting to the novel but uncharted on-device compiler infrastructure.

Our contributions. In this paper, we present a compiler-based solution that can not only be used to study the feasibility of re-instantiating previous solutions such as dynamic, intra-application taint tracking and dynamic permission enforcement, but, moreover, provides a more robust, reliable, and integrated application-layer instrumentation approach than previously possible. Concretely, we make the following contributions in this paper.

Study of the ART compiler suite. Since the novel ART compiler suite, *dex2oat*, is still uncharted, our solution required first and foremost a thorough study of the newly introduced *dex2oat* compiler. We provide, to the best of our knowledge, the first in-depth, comprehensive study of the internals of ART's compiler suite. In particular, we deep-dive into its most recent backend called *Optimizing* that became the default with Android 6 Marshmallow. Those new insights not only allow us to implement compiler-based security solutions, but also form expert knowledge that facilitates independent research on the topic.

Compiler-based app instrumentation. We design and implement a novel approach, called *ARTist* (ART Instrumentation and Security Toolkit), for application instrumentation based on an extended version of ART's on-device compiler *dex2oat*. Our system leverages the compiler's rich optimization framework to safely optimize the newly instrumented application code. The instrumentation process is guided by static analysis that utilizes the compiler's intermediate representation of the app's code as well as its static program information in order to efficiently determine instrumentation targets. A particular benefit of our solution, in contrast to alternative application layer solutions (i.e., bytecode or binary rewriting), is that the application signature is unchanged and therefore Android's signature-based same origin model and its central update utility remain intact. We thoroughly discuss further benefits and drawbacks of security-extended compilers on Android in comparison to bytecode and binary rewriting. Our results provide compelling arguments for preferring compiler-based instrumentation over alternative bytecode or binary rewriting approaches.

Feasibility study for compiler-based taint tracking. To demonstrate the benefits of a solution such as our *ARTist*, we conduct a case study on whether compiler-assisted instrumentation can be utilized to realize a dynamic intra-application taint tracking solution. Our resulting prototype is evaluated using microbenchmarks and its operational capability is shown using an open source test suite with known ground truth.

Outline. The remainder of this paper is structured as follows. In Section 2, we present the results of our study of the *dex2oat* compiler and its *Optimizing* backend. We analyze the requirements for an application-layer instrumentation solution in Section 3 and compare bytecode and binary rewriting with compiler-based approaches. We present our *ARTist* design in Section 4. Section 5 illustrates use cases for *ARTist*, followed by a more detailed case study on compiler-assisted taint tracking. We discuss limitations and future work of our solution in Section 6 and conclude this paper in Section 7.

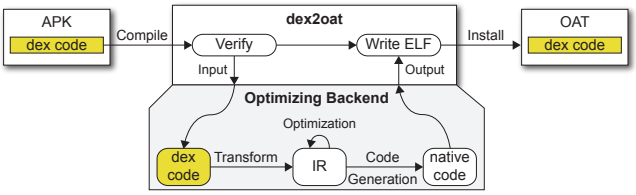

Figure 1: A high-level overview of the *dex2oat* compiler using *Optimizing* backend including the transformation to the IR, optimizations, and native code generation.

2. BACKGROUND ON ART AND *DEX2OAT*

We provide general background information on Android's managed runtime to set the context of our compiler extensions (Section 2.1), and afterwards present technical background information on the compiler suite *dex2oat* (Section 2.2) and in particular on its *Optimizing* backend (Section 2.3).

2.1 Android Runtime

Android is essentially a Linux-based operating system with an extensive middleware software-stack on top of the kernel. The middleware provides native libraries, a feature-rich application framework that implements the Android SDK, and a managed runtime on top of which system as well as third-party applications and a small number of framework services are executed. The runtime executes bytecode generated from Java-based applications and Android's SDK components. The runtime provides the code executed within its environment the necessary hooks to interact with the rest of the system, such as the operating system, the application framework services, or the native Android user space (i.e., components running outside the managed runtime). Every process executing an application runtime environment is usually forked from a warmed-up process, called Zygote, which has all necessary libraries and a skeleton runtime for the app code preloaded.

Runtime prior to Android 5. On Android devices prior to version 5, the runtime consisted of the DEX bytecode interpreter (or Dalvik virtual machine), which was specifically designed for devices with constrained resources (e.g., register-based execution model instead of stack-based). It executes Dalvik executable bytecode (short dex), which is created from the Java bytecode of applications at application build time. Thus, every application package ships the dex bytecode compiled from the application Java sources. Additionally, since Android version 2.2, Dalvik uses just-in-time compilation of hotspot code segments in order to improve the runtime performance of applications.

Runtime since Android 5. With Android 5, Google moved over from an interpreter-based app execution to an on-device, ahead-of-time compilation of apps' dex bytecode to native code that is executed in a newly introduced managed runtime called *ART*. This shift in the runtime model was intended to address the app performance needs of Android's user and developer base. The new compiler suite was designed from scratch to allow for compile time optimizations that improve application performance, start up time, battery lifespan, and also to solve some well-known limitations of the previous interpreter-based runtime, such as the 65k method limit[1]. In particular, Google made the *Optimizing* compiler backend, which was introduced as opt-in feature in Android 5, the default backend in Android 6. In the following Sections 2.2 and 2.3, we will elaborate in more technical details on this new compiler suite and in particular on the *Optimizing* backend.

Prior documentation of ART. Even though the Android source code is publicly available as a part of the Android Open Source Project (AOSP), little attention has yet been given to ART from a security researcher's perspective. Paul Sabanal had an early look [29] at the Android Runtime right after its silent introduction as a developer option on Android 4.4 KitKat. Beside providing information on the ART executable file formats, the paper discusses the idea of hiding rootkits in framework or app code, assuming root access has already been granted. However, especially in its early phase, the Android Runtime has undergone frequent changes, which by now make this documentation unfortunately outdated.[2]

Another work [8] focuses on fuzzing the new runtime with automatically generated input files in order to detect bugs and vulnerabilities. While it provides some high-level overview on the compiler structure and its backends, it unfortunately omits any deeper information on the *Optimizing* backend.

Thus, this background section servers also the purpose of filling a gap in the technical documentation of those new Android features.

2.2 *DEX2OAT* Compiler Suite

Android's on-device compiler *dex2oat* is responsible for the validation of applications and their compilation to native code. It was designed from scratch to be highly flexible and of modular structure, providing numerous configuration possibilities, compiler backends, and native code generators for supported Android platforms. The general workflow of the compiler suite is depicted in Figure 1 and its steps will be explained in the remainder of this section. Providing a full technically documentation of the entire compiler suite and all its intricacies would unfortunately exceed the space limitations of this paper. Therefore, in this section we only focus on the parts relevant to this paper.

2.2.1 *Input File Format*

As an input format, *dex2oat* expects the very same dex files that DVM used to interpret. This strategical decision ensured that neither developers nor app store operators needed to

[1]http://developer.android.com/tools/building/multidex.html

[2]E.g., compare the documented oat version 45 and the current version 63 at the time of this writing.

adapt their code to ART. Developers still upload their apps as Android Application Package (APK) files that bundle the app's code with its resources. When a new app is installed on the device, *dex2oat* compiles the app dex bytecode and the ART runtime executes it, which is completely transparent for the end user. Using this strategy, ART is still compatible with the old Android app base without enforcing a fallback to interpretation, which would loose all benefits that the new compiler provides.

2.2.2 Compilation

Before the actual compilation is performed, each input dex file is checked for validity. Those checks are more extensive and stricter than those implemented in the DVM in order to allow for state-of-the-art code optimizations. The compilation itself is done on a per-method base and can be parallelized. *dex2oat* delegates the actual compilation completely to the backend and only writes the results of the compilation to an oat file along with the original dex code. There are three compilation phases shared between all backends:

Transformation: A graph-based intermediate representation (IR) is created from the dex code. Depending on the actual backend, multiple IRs are possible.

Optimization: Given a populated IR graph, the code is optimized. Each backend provides its own set of optimization measures, ranging from very basic techniques to state-of-the-art algorithms.

Native code generation: The IR nodes are transformed to native code using a code generator for the specific CPU architecture of the current platform. The level of sophistication of the register allocation algorithm and implementation of the code generator depend on the backend.

Backends. On an Android stock device running version 5 (Lollipop) or higher, *dex2oat* can choose between two different backends, *Quick* and *Optimizing*. Although *Quick* was *dex2oat*'s default backend until Android 6, we focus in the remainder of this section and paper on the newer *Optimizing* backend. This choice is not only motivated by the fact that *Optimizing* is the default backend since Android 6 but also by the fact that *Quick* is essentially derived from Dalvik and lacks a sophisticated IR that can support state-of-the-art compiler optimizations—including sophisticated security-oriented algorithms. However, *Optimizing* is designed completely from scratch and little is yet known about its internal structure and design. In Figure 1 the compilation steps of *Optimizing* are depicted. More insights on the inner workings of the new default compiler backend will be provided in Section 2.3.

2.2.3 Oat File Format

Oat files are Android's new file format for apps that are ready to be loaded and executed by the ART runtime. Even though the format was newly created for the Android platform, technically speaking oat files are specialized ELF shared objects that are loaded into processes, i.e., loading a compiled app into an application process is comparable to loading an (ELF) shared library into the process space of a dynamically linked executable. Besides the native code generated with *dex2oat*, oat files contain the complete original dex code,

```
01: public String getID() {
02:
03:   TelephonyManager tm =
04:     getSystemService(TELEPHONY_SERVICE);
05:
06:   String id = tm.getDeviceId();
07:
08:   if(id != null) {
09:      id = prefixID(id);
10:   } else {
11:      id = "N/A";
12:   }
13:
14:   return id;
15: }
16:
17: public String prefixID(String id) {
18:   String prefix = "ID: ";
19:   String result = prefix + id;
20:   Log.d(TAG, prefix + id); // leak id!
21:   return result;
22: }
```

Figure 2: An example code snippet containing a leak of the device's phone number to the logging facility.

which is required to hold up consistency between the code that the developer wrote in Java, the dex code that used to be interpreted, and the compiled code, or to allow fall back to interpretation mode when debugging apps.

2.3 Optimizing Intermediate Representation

We introduce insights into *dex2oat*'s *Optimizing* backend, which we derived mainly from the AOSP source code of the ART project's master branch between June 2015 and February 2016. *Optimizing*'s intermediate representation is essentially a control flow graph on the method level, which the Android developers denote as HGraph. It is further enriched with structural data about the program and populated with instruction nodes, denoted as HInstructions. Figure 2 presents an example Java code and Figure 3 presents the resulting[3] HGraph of the `getID` function in the *Optimizing* IR. We will come back to this example in our case study in Section 5.2.

2.3.1 HGraph

The HGraph serves as the single intermediate representation of the app code. When the graph is created, dex instructions of the app's bytecode are scanned one by one and the corresponding HInstructions are created and interlinked with the current basic block and the graph. In order to allow for complex optimizations, the graph is transformed into a *single static assignment form* (SSA). Pairs of value definitions and usage, so-called *def-use-pairs*, are created during a liveness analysis and explicitly interlinked afterwards. At this point, *phi nodes* are introduced where static analysis cannot reliably decide which value will be assigned at a given position.

In this form, the graph is amenable to a multitude of possible optimizations. The available optimizations includes algorithms such as *BoundsCheckElimination* to remove redundant bounds checks, *GVNOptimization* to remove duplicate code, dead code elimination, or *loop invariant code motion* to

[3]Presented code is simplified and limited to relevant instructions for the sake of readability.

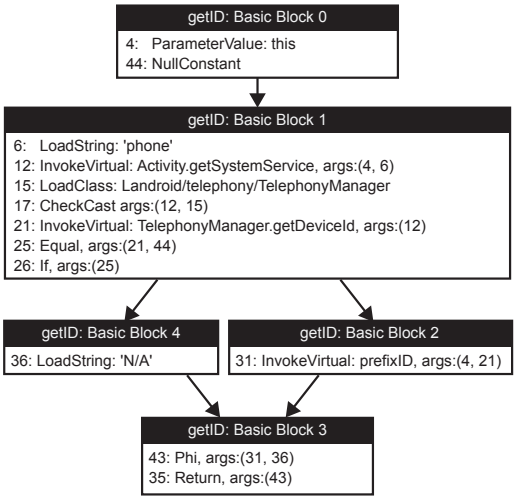

Figure 3: Generated IR in SSA form for the `getID()` method in Figure 2.

optimize hotspot code in loops. In the following Section 4, we will show how this form is also amenable to security-oriented instrumentation, thus supporting compiler-based security solutions on Android, such as dynamic taint tracking (see Section 5.2).

2.3.2 HInstructions

The HGraph nodes roughly correspond to dex instructions. The supplementary material provides an overview of the dex instructions and their corresponding HInstructions counterparts. Beside this transformation, nodes in the HGraph have additional attributes that have no equivalent in dex bytecode (e.g., an SSA index). The HInstructions distinguish between arguments and inputs. While the former corresponds to the arguments given to an operator or method, the latter encodes additional dependencies that may not be immediately observable given only the underlying dex code, as in the case of static method invocations that in addition to their arguments have an HLoadClass or HClinit as their input. All HInstructions share a basic set of information: Type, inputs, uses, id, and further data is attached to each node in order to ease the creation of and working with the HGraph. Each node is uniquely identified within the graph by its id that is assigned and incremented continuously during node creation. The type can be void for methods that have no return value, not for strings and object types of any kind, and additionally any of the Java primitive types. In order to get the actual object type, a fallback to the original dex file is required. This loose coupling between HInstructions and dex instructions as well as the presence of a method local dex program counter in each node show that the IR nodes are not completely independent of the original dex file.

Semantic consistency. In addition to the instructions that represent the original application logic, the HGraph also contains meta-instructions to preserve the semantic consistency between the original Java code of the developer, the dex bytecode shipped with APKs, and the native bytecode actually executed in ART. First, additional instructions are inlined into the graph to support meaningful debugging (e.g., to map from segmentation faults in ART to actual stack traces) and to conduct various forms of runtime checks (e.g., checking type casting, bounds checking, division-by-zero checks, or null pointer exceptions). Second, instructions to represent so-called suspension points are added, which effectively subdivide the application code into multiple chunks. Each suspension point between two chunks acts as a synchronization point between native code and original dex bytecode in the program execution and also serves as an entry point for garbage collectors or debuggers.

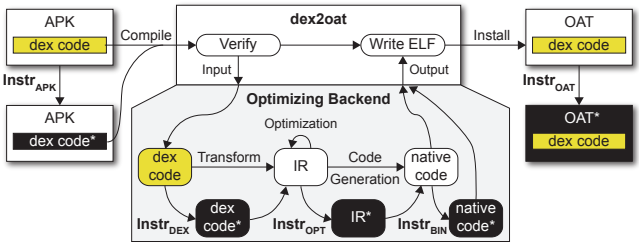

Figure 4: The code instrumentation points before, during, and after the compilation for different representations of the app code. Instrumented code is depicted in black boxes.

3. THE CASE FOR COMPILER-ASSISTED SECURITY ON ANDROID

A dedicated line of work, including the TaintDroid project and its derivatives (such as [22, 2]), relied on instrumentation of the now abandoned Dalvik virtual machine. As a consequence, the research community faces the dilemma on how to continue this line of work and is left with two choices (see Figure 4): Either compensating the missing runtime instrumentation through app rewriting techniques—dex bytecode ($\mathsf{Instr}_{\mathrm{APK}}$) or binary ($\mathsf{Instr}_{\mathrm{OAT}}$)—or taking advantage of Android's new compiler suite ($\mathsf{Instr}_{\mathrm{DEX}}$, $\mathsf{Instr}_{\mathrm{OPT}}$, and $\mathsf{Instr}_{\mathrm{BIN}}$). Although dex bytecode rewriting is well-established in contexts such as inline reference monitoring [23, 16, 15, 7] and taint analysis [31], and ART now supports porting binary rewriting techniques from commodity systems, we build in this paper on compiler-based instrumentation to not only re-instantiate previous approaches that relied on Dalvik VM instrumentation, but also to explore novel security solutions that leverage the compiler features.

In the following, we analyze the concrete requirements that an instrumentation solution should provide and discuss for each of the above approaches (i.e., bytecode rewriting $\mathsf{Instr}_{\mathrm{APK}}$, binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$, and compiler-based instrumentation) their respective benefits and shortcomings in fulfilling those requirements. Table 1 provides a summary of our requirements analysis.

R1. Enforceable security policies. Each of the three approaches operates on one of the different representations of the same app code, i.e., bytecode, IR, or binary. Hence, all three approaches are identical in their capabilities of instrumenting the code and none of the solutions addresses any security policy alone.

R2. Strong security boundary. Both rewriting and the compiler-based approach rely on injecting monitoring code into the app's process space and can therefore not provide a strong security boundary between monitoring and (potentially) malicious app code (✗), e.g., native code. Thus, all of them can only provide security guarantees for at most

	Bytecode rewriting	Compiler-based	Binary rewriting
R1. Enforceable security policies	identical		
R2. Strong security boundary	✗	✗	✗
R3. Application layer only	✓	✓	✓
R4. User privilege only	✓	(✗)	(✗)
R5. Platform independence	✓	✓	✗
R6. Signature preservation	✗	✓	✓
R7. Robustness against optimization	✗	✓	✓
R8. Integrated approach	✗	✓	✗
R9. Supported versions	all	6+	5+

✓= fulfilled; ✗= not fulfilled

Table 1: Comparison of security and deployment features between bytecode rewriting, compiler-based instrumentation, and binary rewriting.

honest-but-curious apps.

R3. Application layer only. All approaches can be implemented purely at application layer (✓). Deploying bytecode rewriting techniques $\mathsf{Instr}_{\mathrm{APK}}$ in form of separate apps has been presented in the literature [23, 16, 15, 7]. A compiler-based solution can be deployed as a separate app that ships and controls the security-instrumented compiler suite (see also Section 4.2). For both, the compiler-based approach and the binary rewriting, the main requirement is access to the storage location of applications' oat files, which does not require system modification.

R4. User privilege only. While dex code is freely available for non-forward locked apps, accessing applications' oat files makes elevated privileges necessary. However, in Section 6.1.3 we discuss an approach that would allow both, binary and compiler-based rewriting, to circumvent this problem without requiring elevated privileges.

R5. Platform independence. Bytecode rewriting $\mathsf{Instr}_{\mathrm{APK}}$ (✓) and compiler-based instrumentation (✓) can be applied on all platforms supported by Android, since they modify the code before platform-dependent native code is generated. Binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$, in contrast, depends on the actual hardware architecture of the platform (✗).

R6. App signature preservation. App signatures are the foundation of Android's same origin model that governs the app update policy or sharing of resources between apps, like a common process or UID. Consequently, modifying bytecode $\mathsf{Instr}_{\mathrm{APK}}$ and the resulting obligation to resign and repackage apps breaks this same origin model (✗). In contrast, compiler-based instrumentation (✓) and binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$ (✓) do not modify the original app package and therefore do not invalidate the signature.

R7. Robustness against code optimization. The instrumentation point determines whether any instrumented code will be subject to optimization at compile time. Applying optimization algorithms to instrumented code has the potential to interfere with the semantics of the modification through, e.g., instruction reordering, inlining, or similar techniques of state-of-the-art compilers like *Optimizing*. Current bytecode rewriting approaches $\mathsf{Instr}_{\mathrm{APK}}$ are applied before compilation, thus any instrumentation has to be robust against optimization by *Optimizing*—an aspect not yet further investigated by contemporary research (✗). On the other hand, binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$ is restricted to instrumenting optimized code (✓), but misses the chance to reuse the rich *Optimizing* framework to also optimize added security code. The sweet spot is compiler-based instrumentation that provides full control over which optimizations are applied when and in which ordering (✓). Even more, this enables creating optimizations that are specifically tailored towards improving the instrumented code by utilizing the static program information that is present in the compiler.

R8. Integration into toolchain. Integrating an instrumentation system into an existing toolchain ensures perpetual development and fixes by the community as well as access to established and well-tested tools and frameworks. In this case, even though the ART project is open source and therefore open to the community, the compiler is mostly maintained by Google itself. Consequently, compiler-driven solutions that do not break with the toolchain's regular functionality, benefit from the continuing improvements (✓). In the case of *ARTist*, the amount of code that needed to be changed is minimal and therefore easy to adapt for newer versions of the toolchain. Bytecode rewriting $\mathsf{Instr}_{\mathrm{APK}}$ and binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$ are developed separately from the toolchain and do not reap those benefits (✗).

R9. Version support. While bytecode instrumentation $\mathsf{Instr}_{\mathrm{APK}}$ can be applied to *all* Android versions, compiler-based approaches and binary rewriting $\mathsf{Instr}_{\mathrm{OAT}}$ depend on ART and therefore can only be applied since Lollipop (*5+*), where a compiler-based solution (as presented here) should utilize the *Optimizing* backend on Android *6+* in preference to *Quick*.

Sweet spot. In conclusion, comparing the security and deployment features that the three available instrumentation approaches provide, a compiler-based approach has very appealing properties and occupies a sweet spot among all approaches.

4. *ARTist* DESIGN

In this section, we present the architecture of the *ART Instrumentation and Security Toolkit*. *ARTist* consists of two separate components: a security-instrumented compiler (*sec-compiler*) and an app to deploy the compiler (*deployment app*). The *sec-compiler* is our implementation of a compile-time instrumentation tool that is based on the *dex2oat* compiler. The latter is a regular Android application that ships, deploys, and manages the *sec-compiler*.

4.1 Security-Instrumented Compiler

Choice of instrumentation point. The general concept of security-instrumented compilers is not restricted in its modifications of the compilation code. Given *dex2oat*'s modular design, there are immediately multiple possibilities apparent where app modifying code could be placed. For instance, *dex2oat*'s design would easily allow porting bytecode and binary rewriting approaches ($\mathsf{Instr}_{\mathrm{DEX}}$ & $\mathsf{Instr}_{\mathrm{BIN}}$) into the

compiler infrastructure (cf. Figure 4). Of the different choices, *ARTist*'s *sec-compiler* is concretely designed to operate on the intermediate representation of *dex2oat*'s *Optimizing* backend ($\mathsf{Instr}_{\mathrm{OPT}}$), where the existing optimization infrastructure and static code information in the *Optimizing* IR allow for efficient and precise code modification. More precisely, our app instrumentation code is realized as an HOptimization, which, as a result, is no different than other optimizations in the sense that they are provided with required information, such as the current method's HGraph, and are modularly integrated into the optimization workflow. As HOptimization, our security instrumentation logic has full control over the ordering and execution of optimizations, which opens up the opportunity to optimize the already instrumented code by creating or applying compatible optimizations that improve the performance of the security code.

Generally, using the HOptimization interface one can extend the compiler with custom functionality that is decoupled from *dex2oat*'s code base. We will refer to those independent extensions as *Modules* for the remainder of this paper.

Spotting instrumentation targets. HGraph supports the visitor pattern [19] that enables us to iterate over, inspect, and modify each single HInstruction of the app's code. In contrast to method hooking techniques, we can therefore operate at the instruction level. In *ARTist*, HGraphVisitors are primarily used to identify instrumentation targets and apply the desired modification. However, they can also be utilized to bootstrap static analysis. We will see concrete implementations using a visitor to collect instrumentation sites for our dynamic permission enforcement system in Section 5.1 and starting points for backward slicing in our taint tracking case study in Section 5.2.

Modification capabilities. Given an instrumentation target in the form of an HInstruction, there are several possibilities for modification like changing types, inputs, or even removing/replacing the instruction altogether. It is also possible to instantiate new instructions of any type and inline them before or after the current target. *ARTist* provides a new API dedicated to automate such modifications if provided with the information which methods, where and what to instrument. Since we are only generating nodes in the form of HInstruction objects and insert them into the HGraph, we do not have to modify the generation of native code from the IR. This means that the code generator is agnostic towards our changes and receives no unexpected instructions or structures. The result of this integrated solution is that we still take advantage of the robustness of *Optimizing*'s code generators, which are well-tested, constantly improved, and in productive use on every stock Android phone running version 6+.

Configuration. Using *Modules*, the instrumentation and modification process is already flexible. To further increase flexibility, *Modules* can, in turn, depend on policy configuration files that govern the instrumentation process. While the design of such policy files highly depends on the concrete *Module*, there are recurring and common patterns, for instance, the amount and type of instrumentation targets that should be detected, as demonstrated in Section 5.1. In general, this allows adaptation of existing instrumentation solutions to new targets or provisioning them with new security policies.

4.2 Compiler Deployment App

Responsibility of the *deployment app* is to deploy the *sec-compiler* at application layer in addition to the system's *dex2oat* binary. Using *deployment app*, one can create security-instrumented versions of installed applications by re-compiling the apps' bytecode with *sec-compiler* and replacing the oat files stored[4] on filesystem. To make the Android runtime agnostic to this instrumentation, two particular challenges had to be overcome. First, Android has mechanisms in place to verify that oat files correspond to their respective apps and that the paths of the oat files are correct. Our implementation solves this challenge by rewriting paths and checksums to match those that the system *dex2oat* would have generated. Second, the oat files are by default stored at and loaded from a protected location to which 3^{rd} party apps have no access. A naïve solution to this problem would be to require extended privileges for our *deployment app* (e.g., a dedicated SELinux type or root on security-relaxed after-market ROMs). We discuss alternatives to the naïve approach in Section 6, which abstain from extended privileges by using app virtualization or reference hijacking. After solving those challenges, the Android default runtime will load the instrumented oat file while remaining agnostic to the fact that we replaced it.

Executing the compiler. Instead of shipping *deployment app* with a statically linked *dex2oat* binary that includes our *ARTist* extensions, we opted for utilizing a copy of Android's default *dex2oat* binary and leveraging its modularity to ship our extensions to the compiler suite as separate libraries. We use the LD_LIBRARY_PATH environment variable to ensure that our *dex2oat* loads and dynamically links our *ARTist* libraries, such as libart-compiler.so, from the assets directory of the *deployment app*.

Inlining custom code. While the instrumentation with *ARTist* already provides powerful tools to modify the application, most security solutions require an additional custom code library within the app (e.g., additional taint tracking logic in Section 5.2). To facilitate adding custom code to an instrumented app, *deployment app* has a preprocessing step that is executed before the app's bytecode is compiled. This step utilizes the DexMerger utility to combine the app's original bytecode with the additional code library. During compilation, connections between original and new code are built in form of invocations of the added code's methods.

5. USE CASES

We demonstrate the applicability and usefulness of our system by discussing several use cases out of which we exemplarily realized two as *ARTist Modules*. First, we implemented an Inline Reference Monitor (IRM) injection *Module* to allow for dynamic permission enforcement. Second, we conduct a case study on realizing intra-app taint tracking through inlining of taint tracking code. In addition, we discuss further ideas for *ARTist Modules*.

[4]Located at `/data/app/<package-name>-1/oat/arm/base.odex`

5.1 IRM for Dynamic Permission Enforcement

In the literature, Inline Reference Monitoring (IRM) is mostly implemented by modifying the bytecode before the installation [23, 16] or by hooking into an application's method at the caller or callee side at runtime [7]. By utilizing a security-instrumented compiler, IRM can be implemented without the need to resign and repackage apps as it is required by established approaches. Moreover, *dex2oat*-based IRM can operate at instruction granularity instead of at the method level. Those capabilities are showcased by our IRM injection *module* that allows for dynamic permission enforcement, as shown by [7, 23, 33] on Android versions before Marshmallow.

The module is split into two distinct parts, the code injection routine that will inline permission enforcement code and the accompanying library that acts as a policy decision point. While the first directs the instrumentation process at installation time, the latter enforces the user's policy at runtime.

Code injection. We first utilize *ARTist* to locate the call sites of permission-protected SDK methods that are defined in a policy configuration file. Afterwards, *ARTist* injects additional calls to our companioning library right before the call sites to check whether the critical method invocations should be allowed. This ensures that the control flow is diverted to our policy decision point before the execution of permission-protected methods.

The limitations imposed by the choice of this rather basic strategy are discussed in Section 6.2.2.

Policy decision point. The library that our *Module* injects into target apps provides methods to check their current state of permissions. Based on the given user permission policy, the library either allows or rejects the execution of a protected SDK method.

5.2 Case study: Taint Tracking

Established approaches for dynamic taint tracking on Android [17] rely on instrumenting the by now scrapped DVM for intra-application taint tracking or directly rewrite bytecode [31]. In this case study, we explore the applicability of a compiler-based instrumentation framework like *ARTist* to re-instantiate intra-app taint tracking for applications on Android version 6 and higher. That is, through a prototypical implementation, we want to investigate whether inlining taint tracking logic into the application code base with *ARTist* at compilation time can be a surrogate for solutions prior to Android version 5. Please note, that this case study does not aim at a full replacement of existing solutions like TaintDroid [17], but at demonstrating a new potential foundation for future taint-tracking on Android.

5.2.1 Module Design

In general, we want to track information as it flows through the code using tracking logic inlined by a new HOptimization in the *Optimizing* backend. However, simply assigning each single value that should be tracked a taint tag and updating the tag for each single instruction operating on it will incur a major performance penalty. To minimalize the runtime impact, we split our approach into two phases: *analysis* and *instrumentation*. In the *analysis* phase, we identify flows of tainted information between *sources* and *sinks*. By restricting ourselves only to those relevant flows of the values we are interested in, we avoid generating irrelevant but costly taint tracking code for parts of the method that never actually influence the data that is obvserved and gain noticable performance improvements over more naïve taint tracking. In the *instrumentation* phase, code will be inlined that creates, propagates, and checks the taint values along the identified data flows. Our combined analysis and instrumentation achieves flow-, path-, object-, and context-sensitive taint tracking.

While [31] and [25] also utilize static analysis to optimize and guide the instrumentation process, both assume a holistic view on the application in form of a control or data flow graph. In contrast, *dex2oat* backends operate on a per-method level, leaving the primary challenge for our taint tracking *Module* to achieve similar tracking properties while inspecting one method at a time. A naïve solution to this problem would be to retrofit the compiler suite to provide an application-wide view and instrumentation. However, our prototype demonstrates how we can still achieve taint tracking for the whole application while restricting ourselves to a per-method view and instrumentation. To this end, we introduce in the following a new design for storing and propagating taint tags, in particular we have to refine the definitions of *sink* and *source*.

5.2.2 Analysis Phase

In order to optimize the instrumentation with taint tracking code, we exploit the processing features (e.g., HGraph's Visitor [19] pattern support) of the *dex2oat* compiler to detect the data flow sources and sinks and afterwards use its static analysis features to identify the relevant data flows and the operations along those flows that have to be instrumented.

Refining source and sink definition. The literature on taint tracking for Android defines sources and sinks as the API methods that input privacy-sensitive information into the application process (e.g., framework functions that return sensitive data, such as the location or telephony API) or, respectively, leak privacy-sensitive information from the application process (e.g., file handles, Internet sockets, or logging facilities). Since *dex2oat* is operating on a per-method level, we cannot assume that our analysis is able to always connect a sink and a source (e.g., when they are located in different methods). To address this problem, we have to connect the data flows of tainted variables across the different methods while maintaining the per-method-based analysis. To this end, we introduce in addition to the above mentioned sinks and sources from the literature—in the following denoted as *global sinks/sources*—new *method-local sinks/-sources*, more precisely HInstructions, which form the entry and exit points for inter-procedural data flows. Thus, global sinks and sources are points of interest for taint tag creation and check, respectively, while local sinks and sources are for inter-procedural tag propagation. For local sinks and sources we differentiate between three categories each: *local sources* include arguments provided to the current method (LSO1), return values from method invocations (LSO2), and values read from fields (LSO3). Conversely, *local sinks* are method invocations that leak values through its parameters from the current method (LSI1), return statements of the current method (LSI2), and field setting instructions (LSI3). At the beginning of the analysis phase we collect all sinks within all methods and in a subsequent step detect all relevant sources for those sinks (see next paragraph).

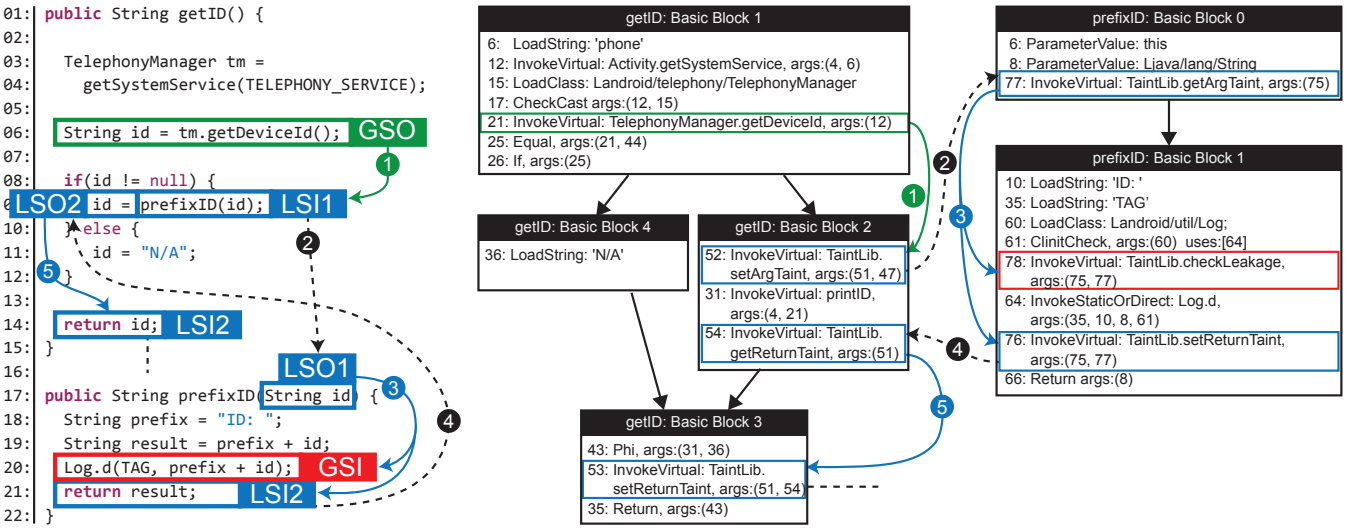

Figure 5: Tracking tainted variable *id* from example of Figure 2. All discovered sinks and sources are marked. Solid lines indicate intra-procedural data flows of tainted variables, dashed lines inter-procedural data flows between local sink-source pairs. Right hand side depicts the inlined taint tracking code to propagate taint tags.

Creating intra-procedural data flows. After collecting all local and global sinks, we create for each sink a backward slice within the currently analyzed method by inspecting the current instruction and recursively tracing back its input parameters depending on the concrete instruction type to discover all sources that influence this sink. For instance, Figure 5 continues the example code from Figure 2. In the Java code on the left hand side, all sinks have been identified (i.e., the parameter `id` passed to function `prefixID` in line 9 is a local sink of type LSI1, the return statements in lines 14 and 21 form local sinks of type LSI2, and in line 20 the call to `Log` forms a global sink). Using backwards slicing (solid lines ①, ③, and ⑤) the local sources in lines 17 (LSO1) and 9 (LSO2) as well as the global source in line 6 (`getDeviceID` call to retrieve device's phone number) have been identified. Each resulting backward slice is defined by its starting point (i.e., the sink) and all found endpoints (i.e., the sources). Constants cannot be tainted and are therefore explicitly omitted as sources. Together those backward slices form the input for the instrumentation phase.

Note that high precision of this analysis is desirable but *not* required for secure taint tracking. Higher precision of this backwards slicing helps in removing irrelevant taint tracking code for the observed values and hence improves performance of the instrumented application's process, but actual taint propagation occurs at *runtime* and the set of our slices contains a superset[5] of the relevant flows for complete tracking of tainted values.

5.2.3 Instrumentation Phase

In the instrumentation phase, we inline code that at runtime creates taint tags for global sources and that checks taints at global sinks. Additionally, we inline code that inter-procedurally propagates taints at runtime from a local sink to a local source, ensuring the data flow of a tainted value across multiple methods correctly propagates the taints.

[5]The use of reflection is an exception to this rule since this poses serious problems for static analyses.

TaintLib. In order to improve the flexibility of our solution by not restricting our system to a specific implementation for storing, updating, and checking taints, we make use of *ARTist*'s modular design and deploy the instrumentation code in form of a new *companion library* called *TaintLib* that is merged by *ARTist* into the app code at compile time. *TaintLib* in turn relies on a policy file that defines the global and local sources/sinks as well as the sources' taints tags. *TaintLib* provides source-type-specific *taint-set* methods, calls to which are inlined at all sources, and sink-type-specific *taint-get* methods, calls to which are inlined at all sinks. By injecting *TaintLib* method calls instead of concrete taint tracking logic, we decouple the instrumentation from the taint management code. For global sources, *taint-get* retrieves and sets the taint tag according to the policy and *taint-set* at global sinks checks[6] the taint tag. In contrast, for local sinks *taint-set* propagates the tag together with the tainted value to the next local source, where it is retrieved with *taint-get*. By instrumenting all methods alike, an implicit contract between all methods is established and fulfilled, i.e., every time a *taint-get* tries to obtain the taint value of a method parameter on the callee side, we assume that the corresponding *taint-set* has been executed in the calling method to provide the taint data. In case the slice contains multiple *sources*, the output of their corresponding *taint-gets* is combined by injecting a call to a combination method that will return the merged taint tag.

To continue our running example, the right hand side of Figure 5 presents the IR of the code snippet with *taint-set* and *taint-get* calls inlined. For instance, the `setArgTaint` call for LSI1 in basic block 2 of `getID` (HInstructions 52) precedes the local sink in HInstructions 31 that invokes the `prefixID` function. The `setArgTaint` instruction transfers the taint of `id` inter-procedurally to the `getArgTaint` instruction in HInstruction 77 of basic block 0 of `prefixID` (dashed line

[6]While a naïve check halts the program when tainted data is about to leak, invoking a sanitizer as suggested by [25] can easily be implemented in *ARTist*.

❷), from where it is intra-procedurally propagated using the backwards slicing information (solid line ❸). Similarly, the taint is propagated back from `prefixID` to `getID` through the return statement and variable assignment (dashed line ❹).

Inter-procedural taint tag propagation channel. In the case of parameters (LSO1 and LSI1) and method returns (LSO2 and LSI2), there are at runtime always pairs of *taint-sets* and *taint-gets*, given by the fact that for each callee method, there is a caller method that also has been instrumented. Combining this with the observation that a caller-callee method pair is always executed by the same thread, the taint propagation can be realized using thread local storage for a taint stack. At the caller side, the taint information is pushed onto a per-thread stack and at the callee side it is popped again, vaguely resembling the x86 calling convention for passing arguments to methods. Keeping in mind that almost every injected *TaintLib* method call accesses the taint information, replacing more straightforward approaches for taint storage (like a single HashMap) with cheaper stack operations also benefits the overall performance of our taint tracking solution.

In the case of field operations (LSO3 and LSI3), we can neither assume them to appear in pairs nor to be executed on the same thread and therefore employ a thread-safe mapping in the form of a ConcurrentHashMap. This, however, raises the challenge of providing easily computable, stable and unique keys. If we consider our taint tags not to store the taint value of a certain value, but of a certain location, we can compute stable identifiers for fields and use them as keys. For static class fields, identifying the specific class and field is sufficient and can be precomputed during compilation. The current implementation injects the computed key as a constant into the HGraph and provides it as an argument to a field *taint-set* or *taint-get*. For object fields, we do not only need to identify classes but concrete objects, which requires runtime information. In this case, we only inject the field identifier as a constant and provide it together with the field's concrete object to a newly added *TaintLib* function. The returned key is robust to object aliasing, so we do not loose track of objects in e.g. collections. Afterwards, we can use this key in a *taint-set* or *taint-get* for the object field.

It is important to note that our approach to taint tracking depends not only on the entity for which we store taints (i.e., variable locations instead of values), but also on the type of data to which we assign taint values. In our model, we track taints only for primitive types and the taint tag of objects is transitively given by their field's tags. In case of non-primitive fields, the rule applies recursively because eventually all objects can be decomposed to primitives. This design decision is motivated by the fact that tracking all *taint-set* and *taint-get* operations on fields and on all method invocations is more fine-grained than storing taint information at the object level.

5.3 Further Use Cases

Dynamic analysis. Compiler-based solutions are inherently well-suited for white box approaches that require an understanding of the application's internals. One example is the taint tracking *Module* described above that re-instantiates TaintDroid-inspired intra-app taint tracking. Other examples are existing works on commodity systems [12, 3] that already utilize compilers for information flow control, which can now be realized on Android as well.

Container solutions. Modifications of the Android runtime environment have been used in the past (for instance *Divide*[7], now part of Google Android for Work) to establish container solutions that, e.g., encrypt file system I/O of apps or restrict inter-application communication. Using a compiler-based approach such as *ARTist*, similar container solutions can be established by replacing the corresponding method invocations (e.g., calls to Java's I/O classes) with calls to injected security-enhanced versions of the same.

Code replacement and compile-time patching. Google has recently started separating security-critical libraries, such as the notorious WebKit, from application packages into stand-alone apps that are called by apps on-demand. This allows Google to maintain those libraries on an ecosystem-wide scale and roll out security patches more effectively. Since *ARTist* is not only able to inject but also to replace or remove code from an app's code base, *ARTist* can also be used to apply *compile-time patches* by replacing vulnerable libraries within apps with fixed versions. In an extreme case, this mechanism could allow for removing entire libraries by mocking all their method invocations (e.g., removing ads), or moving code partitions behind a strong security boundary, such as a dedicated process, and reconnect the code through inter-process communications (e.g., as done in the AdSplit [32] or AdDroid [26] solutions).

Beyond security: profiling and debugging. Besides its application in the security domain, using *ARTist* to inject tracing, debugging or profiling code allows to gain additional insights into third party applications. A basic example is the method call-tracing we employ in our robustness evaluation in Section 6.1.1.

6. DISCUSSION

This section evaluates *ARTist* and its modules in terms of performance, inherent and implementation-specific limitations, and discusses ideas for future work to overcome those and extend our system.

6.1 *ARTist*

We first evaluate and discuss general limitations of the *ART Instrumentation and Security Toolkit*.

6.1.1 Robustness

We briefly evaluate *ARTist* in terms of its robustness by applying an instrumentation routine to a subset of the top apps from different Google play store categories and observing their execution. For the experiment, *ARTist* injects into each method of a target application calls to a carefully crafted tracing method that is merged into the app's code. The tracing method uses stack inspection to determine its caller and prints the corresponding method name to the log. All tests are conducted on a real device (rooted Nexus 5 running Android 6 factory image). Out of 85 non-multidex apps (see Section 6.1.3), 83 apps were successfully instrumented and remained stable when tested (97.64% success rate), indicating the robustness of *ARTist*'s instrumentation capabilities.

[7]http://www.divide.com

Because runtime overhead for instrumented apps largely depends on the concrete code that is injected, we refer to the concrete benchmarks of the *Modules* in Sections 6.2.1 and 6.3.1 for a performance evaluation.

6.1.2 Conceptual Limitations

Native code support. *Optimizing* operates by design on dex input only. Bundled native libraries (i.e., C/C++) that are connected via JNI are never transformed into *Optimizing*'s IR and therefore neither instrumented nor inspected by our prototype. Native code components are a limitation of the attacker model of not only our concept but indeed an open challenge for most of the solutions by Android security research, e.g. code analysis as well as IRM solutions in particular.

Potential fallback to dex. The oat files produced by *dex2oat* still contain the original dex byte code of the app to allow fallback to interpretation mode. Naturally, fallback to interpretation would render our instrumentation of the compiled dex byte code futile. This fallback is currently limited to app debugging, however, no guarantees exist that such a fallback cannot be triggered maliciously. Similarly, dynamically loaded dex code [27, 20] (e.g., via the DexClassLoader) is by default compiled to native bytecode, but no guarantee can be given that dynamically loaded code cannot fall back to interpretation.

6.1.3 Implementation Limitations

Permanence of instrumentation. Instrumentation of an app's oat file might be reverted through an application update or a firmware update where apps are re-compiled. Thus, there exists a window of opportunity for an attacker to start an uninstrumented app after a system or app update. Apps, however, cannot be started programmatically after install/update until the user has started the app manually and both scenarios can be detected by *deployment app* via system notifications (i.e., broadcasts). Assuming that the system notifies the *deployment app* fast enough in order to re-instrument the updated app before the user manually starts the app, the window of opportunity in which an uninstrumented app is started can be closed.

Deployment strategy. In order to create a pure application layer solution, our prototype currently relies on the naïve approach of requesting elevated privileges to replace the installed app oat file with the instrumented version. We can eliminate this requirement by integrating *ARTist* with an application layer only sandboxing solution that provides file system virtualization, such as *Boxify* [6] or *NJAS* [9], or by resetting the execution environment and replacing loaded libraries using reference hijacking [34]. Both approaches enable the manipulation of file paths from the original to the instrumented oat file at application startup time.

Multidex support. Currently our *ARTist* prototype does not support multidex[8] apps. Enabling support for such apps would not only improve coverage of our prototype for more complex apps, but also simplify the merging of additional code by shifting the merge process from the preprocessing step in *deployment app* into the *sec-compiler* by simply providing the additional dex files as input to the compilation process.

[8]http://developer.android.com/tools/building/multidex.html

6.2 Dynamic Permission Module

We briefly evaluate the performance impact of our dynamic permission module and discuss limitations.

6.2.1 Evaluation

The additional security checks inlined by our *Module* are only inserted before permission-protected SDK method calls, so we cannot rely on benchmark apps, because they rarely trigger the added functionality. Therefore, we evaluate the performance impact of our permission checking code using custom microbenchmarks. Table 2 depicts the results of our measurements for calls that are protected by 3 distinct permissions. The overhead encountered in the microbangemarks ranges between 1.18% and 30.65%, thus showing the feasibility of our prototype.

6.2.2 Limitations

Restriction to synchronous calls. In order to demonstrate the straightforward implementation of an *ARTist Module*, we opted for a simple instrumentation strategy that only covers synchronous permission-protected method calls. In result, the current prototype does not support callbacks or asynchronicity and its implementation should therefore be considered a proof-of-concept only.

Best effort permission map. In order to direct *ARTist* to the instrumentation targets, i.e., the application's permission-protected method calls, we utilize a map of methods calls to the permissions enforced by those calls. While the PScout project [5] provides exact API method to permission mappings up to Android version 5.1.1, *ARTist* operates on Marshmallow that requires mappings for Android 6. Consequently, the permission map utilized by our *Module* is a hand picked subset of methods from the v5.1.1 PScout map that did not change for Marshmallow.

6.3 Taint Tracking Module

We evaluate our taint tracking *Module* in terms of feasibility, performance, and the limitations of its current prototypical implementation.

6.3.1 Evaluation

Runtime overhead. We leverage an Android microbenchmark application to evaluate the performance of our prototype. Since our taint-instrumentation only affects the performance of Java code, we specifically chose the *Passmark* benchmark, which does not contain native libraries and implements all benchmarks in Java. Table 3 compares the results of the baseline benchmark with a non-instrumented *Passmark* app to those of an instrumented and taint-aware version. The results show an overhead ranging between 7.74% and 30.73%, which is within an acceptable range for a taint tracking approach that is not fully tuned for performance. This result is also roughly comparable to microbenchmark results of TaintDroid's [17] interpreter-based approach. However, as stated in [24], microbenchmarks are not very representative in user-driven scenarios such as Android apps, so we take this result with a grain of salt.

Overall performance can be enhanced by introducing custom optimizations specifically tailored towards improving

Microbenchmarks				
Tested Method	Permission	Baseline	Instrumented	Penalty
WifiManager.getConfiguredNetworks()	ACCESS_WIFI_STATE	0.681 ms	0.742 ms	8.89%
WifiManager.isWifiEnabled()	ACCESS_WIFI_STATE	0.071 ms	0.072 ms	1.18%
WifiManager.getScanResults()	ACCESS_COARSE_LOCATION	0.452 ms	0.591 ms	30.65%
BluetoothAdapter.startDiscovery()	BLUETOOTH_ADMIN	0.910 ms	0.940 ms	3.32%

Table 2: Microbenchmarks averaged over 60.000 runs. The *baseline* benchmarks measure the pure execution time of the permission-protected call while the *instrumented* benchmarks measure the protected call and the additional permission check.

Passmark			
Test	Baseline	Taint-Aware	Penalty in %
CPU	32521	22526	30.73%
Disk	24893	20777	16.53%
Memory	3627	3346	7.74%

Table 3: Passmark results averaged over 5 runs.

taint tracking code. One approach would be to eliminate *taint-sets* and *taint-gets* that are based on stack operations and cancel each other out, e.g., alternating *pushs* and *pops* of the same tag as seen for methods that return the return value of another method call. Moreover, the analysis phase allows to abstain from instrumenting apps that do not contain any global taint sinks in order to not impact performance at all in this case.

Functional Evaluation. We conducted this case study to research whether intra-application taint tracking can be achieved with an compiler-based instrumentation framework such as *ARTist*, so our functional evaluation focuses on detecting different kinds of data leaks in apps. However, to the best of our knowledge, there is no standardized test suite specifically tailored towards evaluating dynamic taint tracking systems for Android apps, and testing real applications is not feasible because they lack the required ground truth. In order to overcome this unsatisfactory situation, we decided to exploit an open source suite called *DroidBench* [4, 1] that was initially created to benchmark static taint tracking systems. Even though this does not immediately apply to a dynamic system such as ours, we can still leverage the fact that it provides us with an assortment of applications with different but well-defined leakage behavior. Table 4 summarizes our *Module's* results for those tests and categories within scope. Tests for implicit flows, inter-component communication, and reflection are omitted because they currently exceed the scope of our proof-of-concept taint tracking. As we are *abusing* the benchmark suite, we need to be careful which conclusions we draw from the test results. The first insight we gain, however, is that our case study succeeded in showing that intra-app taint tracking can be implemented as a pure application layer solution using compiler-driven instrumentation. The second insight we derive is that, as indicated by lower results such as those for the *Android Specifics* category, our proof-of-concept does not yet catch up with previous works such as TaintDroid. Nonetheless, our work not only shows the feasibility of the approach but also lays the foundation for creating a full-fledged taint tracking system for Android versions above Marshmallow that utilizes compiler-based instrumentation and does not require modification of the operating system.

DroidBench		
Category	Successful Tests	Ratio
Callbacks	14/15	93%
Lifecycle	13/14	92.9%
General Java	14/20	70%
Aliasing	1/1	100%
Android Specifics	5/9	55.6%
Field & Object Sensitivity	7/7	100%
Overall	54/66	81.8%

Table 4: Results for the DroidBench taint tracking evaluation. Broken tests and categories not applicable to our system are omitted.

6.3.2 Limitations

No tracking of implicit flows. Like TaintDroid [17], our system currently does not track implicit flows (i.e., data leakage using control flow dependencies) and malevolent apps could exfiltrate data in a way that is unnoticeable by our prototype. As the TaintDroid authors discuss, mitigating leakage through control flows would require static analysis and access to the app's source code—both of which TaintDroid could not provide. *ARTist* however is already provided with the full app code and it would be highly interesting future work to investigate to which extent the structural program information of the IR and analytical features of the compiler backend (e.g., *Optimizing*) can help to remedy the limitations of customary taint tracking solutions on Android.

Taint tracking boundaries. The compiler is restricted to the app's code base, which introduces imprecision when leaking information through SDK methods, where a *taint-set* at the caller side (developer code) but not the *taint-get* at the callee side (SDK) can be inlined. In particular, and in contrast to object types, storing primitives or strings in collections or sharing them across threads are corner cases where the taints will not be propagated appropriately. This shortcoming can be solved by using pre-computed control-flow models for framework methods [11] to generate corresponding *taint-set* and *taint-get* pairs that model the transition of data through the framework. A preferable technical solution in the future, which removes the potential over-approximations of SDK internal states in control-flow models [11] and which could be of interest beyond taint tracking, is the instrumentation of the *core image*. The core image is a pre-compiled oat file of the framework classes that is pre-loaded into every application process via Zygote. Since the core image is created with *dex2oat* during the device startup once after each system update, it can be instrumented using a *sec-compiler* as in *ARTist*. However, in either case and as in the original work [17], data that already left the phone (e.g., through a network socket) cannot be tracked.

Inter-application communication. Our prototype is currently limited to *intra*-application tracking and lacks support for *inter*-application tracking, for instance, through the file system or Binder IPC. This opens the possibility of confused deputy [18, 14] or collusion attacks [30, 10] to exfiltrate data. Assuming that all installed apps are instrumented, a fix to this problem would be the instrumentation of the I/O method calls in order to write out taints together with the data (e.g., into a file or Binder Parcel) and restore the taints at receiver side. When abandoning the requirement for a pure application-layer solution, our system could also be complemented with the original TaintDroid file system and IPC infrastructure, which is unaffected by the loss of DVM, in order to track taints across applications.

7. CONCLUSION

In this paper, we presented *ARTist*, a *dex2oat* compiler-based instrumentation solution for Android applications that operates at application-layer only. In order to be able to design and implement *ARTist*, we first and foremost had to thoroughly study the yet uncharted internals of the new compiler suite and in particular of its *Optimizing* backend. A deeper understanding of this compiler suite and the new ART runtime is insofar of interest for the security community, since ART and *dex2oat* replaced the interpreter-based runtime (DVM) of Android versions prior to Lollipop (i.e., version 5) and hence also voided applicability of any security solution that relies on interpreter instrumentation (e.g., TaintDroid [17] and its derivatives [22, 2]). We study feasibility of our approach through implementing two distinct use cases. Furthermore our case study highlights the capability of a compiler-based instrumentation framework to re-instantiate basic taint tracking for Android apps on the application layer. In general, our results provide compelling arguments, such as higher robustness and better integration, for preferring compiler-based instrumentation over alternative bytecode or binary rewriting approaches.

8. REFERENCES

[1] Online: https://github.com/secure-software-engineering/DroidBench.

[2] Google code: Droidbox. https://code.google.com/p/droidbox/.

[3] F. Araujo, W. Kevin, et al. Compiler-instrumented, dynamic secret-redaction of legacy processes for attacker deception. In *USENIX SEC'15*.

[4] S. Arzt, S. Rasthofer, C. Fritz, E. Bodden, A. Bartel, J. Klein, Y. Le Traon, D. Octeau, and P. McDaniel. FlowDroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. In *ACM PLDI '14*.

[5] K. W. Y. Au, Y. F. Zhou, Z. Huang, and D. Lie. Pscout: analyzing the android permission specification. In *ACM CCS'12*.

[6] M. Backes, S. Bugiel, C. Hammer, O. Schranz, and P. von Styp-Rekowsky. Boxify: Full-fledged app sandboxing for stock android. In *USENIX SEC'15*.

[7] M. Backes, S. Gerling, C. Hammer, M. Maffei, and P. von Styp-Rekowsky. AppGuard - enforcing user requirements on Android apps. In *TACAS'13*.

[8] A. Bechtsoudis. Fuzzing Objects d'ART: Digging Into the New Android L Runtime Internals. Online: http://census-labs.com/media/Fuzzing_Objects_d_ART_hitbsecconf2015ams_WP.pdf, 2015.

[9] A. Bianchi, Y. Fratantonio, C. Kruegel, and G. Vigna. NJAS: Sandboxing unmodified applications in non-rooted devices running stock android. In *ACM CCS SPSM'15*.

[10] S. Bugiel, L. Davi, A. Dmitrienko, T. Fischer, A.-R. Sadeghi, and B. Shastry. Towards Taming Privilege-Escalation Attacks on Android. In *NDSS'12*.

[11] Y. Cao, Y. Fratantonio, A. Bianchi, M. Egele, C. Kruegel, G. Vigna, and Y. Chen. EdgeMiner: Automatically detecting implicit control flow transitions through the android framework. In *NDSS'15*.

[12] W. Chang, B. Streiff, and C. Lin. Efficient and extensible security enforcement using dynamic data flow analysis. In *ACM CCS'08*.

[13] L. P. Cox, P. Gilbert, G. Lawler, V. Pistol, A. Razeen, B. Wu, and S. Cheemalapati. Spandex: Secure password tracking for android. In *USENIX SEC'14*.

[14] L. Davi, A. Dmitrienko, A.-R. Sadeghi, and M. Winandy. Privilege escalation attacks on android. In *ISC'10*.

[15] B. Davis and H. Chen. Retroskeleton: Retrofitting android apps. In *ACM MobiSys'13*.

[16] B. Davis, B. Sanders, A. Khodaverdian, and H. Chen. I-ARM-Droid: A Rewriting Framework for In-App Reference Monitors for Android Applications. In *IEEE CS MoST'12*.

[17] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. TaintDroid: An information-flow tracking system for realtime privacy monitoring on smartphones. In *USENIX OSDI'10*.

[18] A. P. Felt, H. J. Wang, A. Moshchuk, S. Hanna, and E. Chin. Permission re-delegation: Attacks and defenses. In *USENIX SEC'11*.

[19] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-oriented Software*. Addison-Wesley Longman Publishing Co., Inc., 1995.

[20] M. Grace, W. Zhou, X. Jiang, and A.-R. Sadeghi. Unsafe exposure analysis of mobile in-app advertisements. In *ACM WISEC'12*.

[21] H. Hao, V. Singh, and W. Du. On the Effectiveness of API-level Access Control Using Bytecode Rewriting in Android. In *ACM ASIACCS'13*.

[22] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. These aren't the droids you're looking for: Retrofitting android to protect data from imperious applications. In *ACM CCS'11*.

[23] J. Jeon, K. K. Micinski, J. A. Vaughan, A. Fogel, N. Reddy, J. S. Foster, and T. Millstein. Dr. Android and Mr. Hide: Fine-grained Permissions in Android Applications. In *ACM CCS SPSM'12*.

[24] B. Livshits. Dynamic taint tracking in managed runtimes. Technical report, Microsoft Research.

[25] B. Livshits and S. Chong. Towards fully automatic placement of security sanitizers and declassifiers. In *ACM SIGPLAN'13*.

[26] P. Pearce, A. Porter Felt, G. Nunez, and D. Wagner. AdDroid: Privilege separation for applications and

advertisers in Android. In *ACM ASIACCS'12*.

[27] S. Poeplau, Y. Fratantonio, A. Bianchi, C. Kruegel, and G. Vigna. Execute this! analyzing unsafe and malicious dynamic code loading in android applications. In *NDSS'14*.

[28] G. Russello, M. Conti, B. Crispo, and E. Fernandes. MOSES: supporting operation modes on smartphones. In *ACM SACMAT'12*.

[29] P. Sabanal. Hiding behind ART. Online: https://www.blackhat.com/docs/asia-15/materials/asia-15-Sabanal-Hiding-Behind-ART-wp.pdf, 2015.

[30] R. Schlegel, K. Zhang, X. Zhou, M. Intwala, A. Kapadia, and X. Wang. Soundcomber: A stealthy and context-aware sound trojan for smartphones. In *NDSS'11*.

[31] J. Schütte, D. Titze, and J. De Fuentes. AppCaulk: Data leak prevention by injecting targeted taint tracking into android apps. In *TrustCom14*.

[32] S. Shekhar, M. Dietz, and D. S. Wallach. Adsplit: Separating smartphone advertising from applications. In *USENIX SEC'12*.

[33] R. Xu, H. Saïdi, and R. Anderson. Aurasium – Practical Policy Enforcement for Android Applications. In *USENIX SEC'12*.

[34] W. You, B. Liang, W. Shi, S. Zhu, P. Wang, S. Xie, and X. Zhang. Reference hijacking: Patching, protecting and analyzing on unmodified and non-rooted android devices. In *ICSE'16*.